Centre de recherche sur les transports (C.R.T.)
Centre for Research on Transportation

CRT-2003-42

DECOMPOSITION METHODS AND ITERATIVE COMBINATORIAL AUCTIONS

by

Jawad Abrache
Teodor Gabriel Crainic
Michel Gendreau

December 2003

Université de Montréal
C.P. 6128, succursale Centre-ville, Montréal QC H3C 3J7 Canada
Téléphone: (514) 343-7575 / Télécopieur: (514) 343-7121

crt@crt.umontreal.ca
www.crt.umontreal.ca

HEC ÉCOLE POLYTECHNIQUE MONTRÉAL Université de Montréal

104-02696/

PUBLICATION
CRT-2003-42

DECOMPOSITION METHODS AND ITERATIVE COMBINATORIAL AUCTIONS

by

Jawad Abrache[1]
Teodor Gabriel Crainic[2]
Michel Gendreau[1]

The research was supported in part by the Bell University Laboratories, the Natural Sciences and Engineering Research Council of Canada (NSERC), and by the Fonds québécois de la recherche sur la nature et les technologies (FQRNT).

1. Département d'informatique et de recherche opérationnelle, Université de Montréal, and Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7 (jawad; michelg@crt.umontreal.ca)

2. Département de management et technologie, Université du Québec à Montréal, et Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7 (theo@crt.umontreal.ca)

Département d'informatique et de recherche opérationnelle - Publication #1245

Dépôt légal – Bibliothèque nationale du Québec
Bibliothèque nationale du Canada, 2003

Abstract

Iterative auctions are motivated by the fact that the market-maker often lacks complete and truthful information about the bidders' private valuations of the resources on sale. The literature on the design of iterative mechanisms for combinatorial auctions has addressed only the most basic cases and has been dominated by primal-dual approaches. In this paper, we consider a general production/consumption exchange of interdependent goods, for which we investigate iterative auction mechanisms based on mathematical programming decomposition methods. We focus specifically on two well-known approaches: Lagrangian relaxation and Dantzig-Wolfe decomposition.

Keywords: combinatorial auctions, iterative auction mechanisms, primal-dual approaches, mathematical programming decomposition methods.

Résumé

Le besoin de mécanismes d'enchère itératifs, ou multirondes, est essentiellement motivé par le fait que l'encanteur n'a souvent pas accès aux préférences privées des participants dans le marché. L'état de l'art dans le design de mécanismes d'enchère multirondes indique que l'effort de recherche s'est limité aux cas de figures les plus simples, et s'est principalement inspiré de l'approche primale-duale en programmation linéaire. Nous considérons dans cet article une économie constituée de producteurs et de consommateurs de biens interdépendants, pour laquelle nous élaborons des processus d'enchère multirondes basés sur les méthodes de décomposition en programmation mathématique. Nous mettons l'emphase sur deux approches bien connues : la relaxation lagrangienne et la décomposition de Dantzig-Wolfe.

Mots-clés : enchères combinatoires, mécanismes d'enchère itératifs, approche primale-duale, approches de décomposition.

1 Introduction

The spectacular growth of electronic commerce and information technology has made the Internet a place of choice for conducting business. An immediate consequence of this trend is the creation of many public and private *electronic marketplaces* where huge volumes of goods and services are exchanged every day. The apparition of these new entities has also brought forth many fundamental issues related to their structure and organization. Among these, the question of how the economic rules that govern the market should be designed is of prime importance and contributes to the renewed interest in the economic, game-theoretic, and computational aspects of market mechanism design.

A market mechanism can be defined, in informal terms, as a set of deterministic rules that specify an allocation of the items traded in the market to the participants, as well as the corresponding payments the latter should make or receive. According to the economic theory of mechanism design (Mas-Colell, Whinston, and Green 1995), one distinguishes basically two classes of market mechanisms based on the information exchanged between the market-maker and the participants, *direct-revelation* and *indirect-revelation* mechanisms. The former require the participants to completely and truthfully report their *types* to the market-maker; that is, in the context of a production-consumption economy, for instance, sellers should disclose their production technologies and cost functions and buyers their preferences for the goods and the consumption constraints. Assuming such information is available, the market-maker could then formulate an optimization model that clears the market such that an objective of "effectiveness" (that could be the maximization of the total social welfare of the participants, the revenue of a seller, etc.) is achieved.

One could argue, however, that complete and truthful information revelation is too strong an assumption in most market contexts. First and foremost, participants are generally self-interested agents, which pursue their own goals independently of the market objectives. They may thus be unwilling to disclose, without careful screening, their private data. In many cases, this data may simply be too complex to assess and to communicate in its entirety (Nisan 2001). Finally, a participant may possess only uncertain valuations of the goods, which it hopes to adjust by observing how other participants behave (e.g., auctions of art objects or rare items). Indirect-revelation mechanisms help circumvent some of these problems by departing from the assumption made of an omniscient market-maker with systematic access to all the information. Iterative auctions constitute an important class of indirect-revelation mechanisms that allow for progressive and considerably less forceful disclosure of information. In each round of an iterative auction, participants submit *bids* to sell or buy items, to which the market-maker responds by determining provisional allocations and payments, and by sending "signals" about the state of the auction that give participants impetus to commit further. It is noteworthy that bids neither have to represent complete preferences - they only reflect participants' needs *given the observed signals*, nor to convey truthful information - unless the auction incentives prevents preference misrepresentation.

The design of iterative auction mechanisms for the optimization of distributed systems has

attracted much attention recently. In particular, the Combinatorial Allocation Problem (CAP), which consists in determining a socially-efficient allocation of multiple indivisible goods to several potential buyers, has been thoroughly investigated. For the CAP, Parkes' doctoral thesis (Parkes 2001) surveys the research path toward incentive-compatible, multi-round auctions that reconcile the market's welfare maximization objective with self-interested buyers aiming at maximizing their profits. Most of these auctions can be interpreted as primal-dual algorithms that capitalize on linear programming formulations of the problem.

Mathematical programming decomposition approaches (Geoffrion 1970) offer an attracting alternative. These methods have been used for decades to address large-scale structured optimization problems. Yet, their potential for decentralized decision making, if thoroughly analyzed, has seldom been exploited in the design of auction mechanisms. Hence, this paper aspires to contribute to a better understanding of the market properties of two important decomposition approaches, that is, Lagrangian relaxation and Dantzig-Wolfe (DW) decomposition.

We consider a general combinatorial exchange economy in which self-interested agents trade several different divisible goods. We first formulate the allocation and payment rules of an ideal direct-revelation mechanism for this market, with the objective of maximizing the overall social efficiency of the market. We then show that, under appropriate assumptions, the application of Lagrangian relaxation and Dantzig-Wolfe decomposition to the allocation problem in the centralized market leads to indirect mechanisms in which the agents' own interests are ultimately reconciled with the "global" objective of the market. We futhermore establish that subtle but important differences exist between relaxation and DW-based auction processes, in their use of the information disclosed by the agents through their bids and the assumptions these mechanisms make on the agents' strategic behavior.

The paper makes several important theoretical and practical contributions. While mathematical programming decomposition methods have been presented in the past as market mechanisms (notably Lagrangian relaxation, se de Vries and Vohra, 2003), the analysis has been limited to the one-sided case (that is, the market-maker selling several different items to many buyers). To the best of our knowledge, this is the first attempt to analyze these methods in a many-to-many combinatorial exchange context, in which many sellers and buyers interact. Moreover, the paper presents an original comparison of these methods from informational and strategic perspectives. Finally, the numerical results provide interesting insights into the potential and limitations of auction mechanisms based on the two methods.

The paper is organized as follows. In Section 2, we survey the existing literature on iterative auction mechanisms for combinatorial markets. We put the emphasis on classical tâtonnement processes and auction mechanisms based on the primal/dual approach. We formulate in Section 3 the problems of determining a socially efficient allocation and the corresponding equilibrium prices in a centralized many-to-many direct-revelation market. In the following two sections, we present two relaxation-based methods (using the subgradient algorithm and a bundle method, respectively) and the Dantzig-Wolfe decomposition scheme, and interpret them as iterative auction processes. Finally, we devote Section 6 to an

experimental study that performs a preliminary evaluation of the corresponding auctions.

2 Prior work

The theory of general competitive equilibrium (Mas-Colell, Whinston, and Green 1995) institutes the framework for the econometric analysis of markets involving the trade of several interdependent divisible goods. General equilibrium is specifically concerned with the study of the social efficiency of the allocation of the goods in the market, as well as the determination of states of equilibrium from which the participants would not depart. In particular, *Walrasian* equilibrium, defined in terms of payments based on single-item prices, has been extensively studied in the economic literature and a focal result establishes their existence under certain conditions (Arrow and Debreu 1954).

The price-tâtonnement process of Walras (1874) is probably the first iterative scheme aiming at the determination of competitive equilibria. Walrasian tâtonnement proceeds by adjusting prices of the various goods one by one until equilibration of supply and demand of all the goods is realized, which is only guaranteed if additional conditions that insure the equilibrium *stability* are satisfied (for instance, the *gross-substitutes* condition, which states that the demand for a given item does not decrease if prices of other items increase, is sufficient for equilibrium stability, Arrow and Hahn 1971). Price-tâtonnement has been recently implemented in the WALRAS algorithm (Cheng and Wellman 1998). More generally, it has given rise to computational paradigms such as *market-oriented programming* (Wellman 1996), which model and implement resource allocation problems as distributed systems of autonomous agents reacting to a pricing system. Market-oriented programming approaches have been applied to many areas, including the allocation of transportation services (Wellman 1993), quality-of-service allocation in multimedia applications (Yamaki, Wellman, and Ishida 1996), vehicle routing (Sandholm 1993), power load management (Ygge 1998), and decentralized scheduling (Wellman *et al.* 2001).

Resource-directed quantity-tâtonnement processes form the other important class of market mechanisms based on tâtonnement. In an iteration of a quantity-tâtonnement process, the market-maker determines provisional allocations of goods and participants submit their *marginal* utilities, that is, the utilities for producing or consuming one additional unit of each good. The allocations are adjusted by the market-maker such that participants with higher (lesser) marginal utilities are given more (less) of the goods, until a state in which everybody is ready to pay (receive) the same marginal prices. Simple fixed and variable-stepsize adjustment methods has been suggested by Kurose and Simha (1989), and a more efficient Newton-Raphson search method by Ygge and Akkermans (2000).

The auction algorithm of Bertsekas (1988) for the assignment problem is one of the first and most successful attempts to use an iterative distributed process for solving optimization problems. The algorithm develops basically a multi-round auction and capitalizes on

the economic interpretation of complementary slackness conditions in linear programming as equilibrium conditions for self-interested agents. Hence, in a given iteration of the algorithm, each agent chooses the most "interesting" object at the current prices and bids a new value that beats the current price by a certain increment (computed to maintain complementary slackness throughout the auction). The market-maker then simply allocates the objects to the respective highest bids. Extensions of the auction algorithm to other network flow problems, notably the shortest path and the minimum cost flow problems, have afterwords been presented (Bertsekas 1991, 1992). Another important iterative process for the assignment problem is due to Demange, Gale, and Sotomayor (1986). The process is a price-ascending auction in which prices on over-demanded goods are raised across the auction rounds, and can be seen as a particular implementation of the primal-dual algorithm (Bikhchandani *et al.* 2001). More importantly, the authors show that by choosing *minimal* over-demanded sets of items, the auctioneer produces "second-price" Vickrey payments. Demange, Gale, and Sotomayor mechanism is thus an incentive-compatible iterative auction for the assignment problem.

The study of iterative auctions for the Combinatorial Auction Problem (CAP) has been at the heart of recent research. The CAP is the problem of determining welfare-maximizing allocations of several indivisible items to buyers that have preferences for bundles of items. Bikhchandani and Mamer (1997) investigate the existence of Walrasian equilibrium prices for the CAP and establish that it relies on the integrality of the LP relaxation of the CAP. For their part, Bikhchandani and Ostroy (2000) suggest extended formulations of the CAP that imply nonlinear - and possibly discriminatory - bundle equilibrium prices always exist and correspond to optimal dual solutions of the associated formulations. This allowed the development of primal-dual implementations (e.g., the iBundle family of iterative auctions, Parkes 2000). The existence of LP formulations of the CAP leads to iterative auctions implementing Vickrey payments in particular cases. Thus, when an *"agents-are-substitutes"* condition holds, which states that the impact on the optimal allocation of a group of participants exceeds the sum of the individual impacts of the participants, Bikhchandani and Ostroy (2000) show that minimal competitive equilibrium prices correspond to Vickrey payments. Under the additional assumption of *gross substitutes*, primal-dual implementations have been suggested by Gul and Stacchetti (2000) and Ausubel (2002).

Decentralized market-based mechanisms have also been suggested for coordinating the flow of goods, information, or services within organizations. Important methodological developments in the analysis of information and decision distribution within an organization include Malone, Yates, and Benjamin (1987), which compares empirically decentralized decision mechanisms (more specifically markets) and centrally-coordinated one (hierarchies), and predicts that advances in IT will induce firms to shift from hierarchies to market-centric processes, by making information more accessible and less costly. Tan and Harker (1999) focus on operation scheduling, for which they quantify centralized and decentralized workflow structures with respect to production, coordination and disruption costs. The results of their study indicate that a distributed design for operation scheduling is more attractive than scheduling by a central authority when job processing times are relatively long, communication costs are low, and agents are reliable. A distributed scheduling framework, based on the

Lagrangian relaxation technique, appears in Kutanoglu and Wu (1999). The authors develop an iterative combinatorial auction scheme - they call *adaptative tâtonnement* - in which jobs requiring processing time submit utility-maximizing bids on time slots on the corresponding machines and a coordinating agent allocates time slots and updates prices in proportion to the excess demand arising for time slots. The authors also show that Lagrangian relaxation using the subgradient algorithm corresponds to a special case of adaptative tâtonnement. Lagrangian relaxation has been similarly suggested by Kim *et al.* (2002) in the design of an electronic brokering process for truckload freight.

3 Centralized market-clearing for a combinatorial exchange economy

Consider an economy with a set of divisible goods on sale and two categories of agents: sellers and buyers. Sellers have the capacity to produce the goods according to their own technology and production cost functions. Buyers, for their part, consume goods either directly or as inputs to a transformation process. Hence, they have preferences for bundles of goods on sale and may also face technological requirements that constrain their consumption.

Let us introduce the following notation:

- $\mathcal{L}$: set of goods;
- $\mathcal{S}$: set of sellers;
- $\mathcal{J}$: set of buyers;
- $q_{s,l}$: quantity of good $l, l \in \mathcal{L}$ produced by seller s, $s \in \mathcal{S}$;
- $q_{j,l}$: quantity of good $l, l \in \mathcal{L}$ consumed by buyer j, $j \in \mathcal{J}$;
- $\mathcal{D}_s$: production feasibility set of seller s, containing all admissible quantity vectors $q_s = \{q_{s,l}\}_{l \in \mathcal{L}}$ that seller s may produce;
- $\mathcal{D}_j$: consumption feasibility set of buyer j, containing all admissible quantity vectors $q_j = \{q_{j,l}\}_{l \in \mathcal{L}}$ that buyer j may consume;
- $C_s(.)$: production cost function of seller s, $s \in \mathcal{S}$: that is, $C_s(q_s)$ is the cost to seller s of producing q_s;
- $V_j(.)$: valuation function of buyer j, $j \in \mathcal{J}$: similarly, $V_j(q_j)$ is buyer j's preference for consuming q_j.

We make the following assumptions on the feasibility sets and the cost and valuation functions:

(A1) Production cost functions are *continuous*, *convex*, *monotone increasing* (i.e., $C_s(q_s^{(1)}) \geq C_s(q_s^{(2)})$, $\forall q_s^{(1)}, q_s^{(2)} \in \mathcal{D}_s$ s.t. $q_s^{(1)} \geq q_s^{(2)}$. Similarly, buyer valuation functions are *continuous*, *concave*, *monotone increasing*.

(A2) Production and consumption feasibility sets are *convex* and *bounded*.

A direct-revelation market mechanism (see Figure 1) would be in this case a one-shot market-clearing process aimed at matching the supply and demand for the various goods. The sellers and buyers need to communicate to the market-maker their production and consumption feasibility sets and their cost and valuation functions, respectively. The mechanism's output is an allocation of the goods and payments the sellers need to make to the buyers.

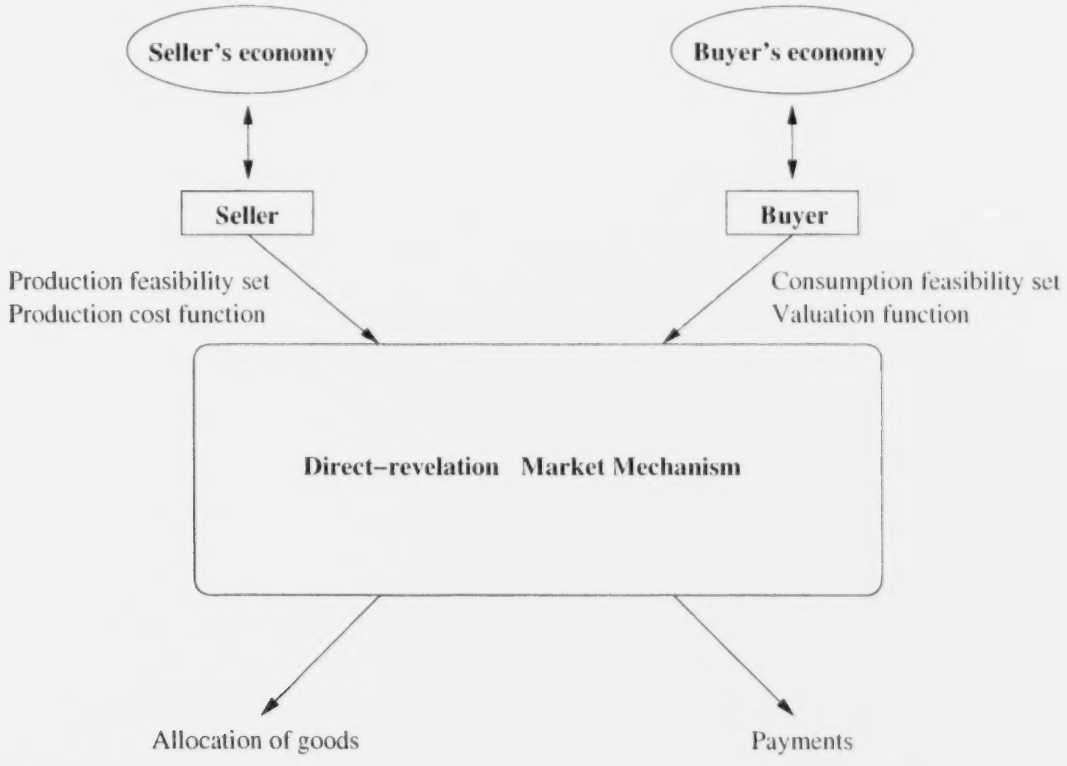

Figure 1: A direct-revelation market mechanism.

The primary goal of the market-maker is to determine a *socially-efficient* allocation, that is, a feasible allocation of the goods that maximizes the global welfare of the sellers and buyers. With the notation above, a socially-efficient allocation is a solution of model (MC):

$$\max \quad \sum_{j\in\mathcal{J}} V_j(q_j) - \sum_{s\in\mathcal{S}} C_s(q_s) \tag{1}$$

$$s.t. \quad \sum_{j\in\mathcal{J}} q_{j,l} - \sum_{s\in\mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \tag{2}$$

$$q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \tag{3}$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \tag{4}$$

The objective of model (MC) is to maximize the market surplus, computed as the difference between the buyers' valuations and the sellers' production costs. Equations (2) matches the demand with the supply, while relations (3) and (4) constrain the quantities traded to be within admissible production and consumption limits.

The concept of pricing equilibria serves the market-maker's second important purpose, which is to reconcile the global objective of the market - social efficiency - with the interests of the sellers and buyers. Prior to the definition of a Walrasian equilibrium, it is convenient to make the following classical assumptions:

(A3) Sellers and buyers have *quasi-linear* utilities, in which money is considered as an exogenous resource. That is, the utility of a buyer j with a wealth level m for consuming q_j is $u_j(q_j, m) = V_j(q_j) + m$ and that of a seller s producing q_s is $u_s(q_s, m) = -C_s(q_s) + m$.

(A4) Sellers and buyers are *price-takers*, i.e., they take market prices as given and ignore the effect of their own actions on prices.

Definition 1 *(Walrasian Equilibrium) The price vector $p = \{p_l\}_{l\in\mathcal{L}}$ constitutes a set of Walrasian prices if there is a feasible allocation $\tilde{q} = [\{\tilde{q}_j\}_{j\in\mathcal{J}}; \{\tilde{q}_s\}_{s\in\mathcal{S}}]$ such that*

1. $V_j(\tilde{q}_j) - p.\tilde{q}_j = \max_{q_j\in\mathcal{D}_j}(V_j(q_j) - p.q_j)$, $\forall j \in \mathcal{J}$; *and*
2. $p.\tilde{q}_s - C_s(\tilde{q}_s) = \max_{q_s\in\mathcal{D}_s}(p.q_s - C_s(q_s))$, $\forall s \in \mathcal{S}$.

4 Market-clearing based on Lagrangian relaxation

The particular structure of the centralized market-clearing formulation (MC) naturally suggests Lagrangian relaxation as an approach to decompose the problem. Hence, consider (MC) as the primal problem and dualize it with respect to the matching constraints (2). Let

$\lambda = \{\lambda_l\}_{l \in \mathcal{L}}$ be the vector of Lagrangian multipliers associated with (2). The corresponding Lagrangian can be defined as:

$$\begin{aligned} L(q;\lambda) \quad &= \sum_{j \in \mathcal{J}} V_j(q_j) - \sum_{s \in \mathcal{S}} C_s(q_s) + \sum_{l \in \mathcal{L}} \lambda_l (\sum_{s \in \mathcal{S}} q_{s,l} - \sum_{j \in \mathcal{J}} q_{j,l}); \\ &\forall q : q_j \in \mathcal{D}_j, j \in \mathcal{J}; q_s \in \mathcal{D}_s, s \in \mathcal{S}; \forall \lambda \in \mathbb{R}^{|\mathcal{L}|} \end{aligned} \tag{5}$$

Consider the dual function $\Theta(\lambda) = \max_q \{L(q;\lambda) : q_j \in \mathcal{D}_j, j \in \mathcal{J}, q_s \in \mathcal{D}_s, s \in \mathcal{S}\}$. The Lagrangian dual problem (LD) is:

$$\min_{\lambda} \Theta(\lambda) \tag{6}$$

It is noteworthy that the Lagrangian dual function can be written as the following problem, $(\mathrm{LR}_{(\lambda)})$:

$$\max_{q} \quad \sum_{j \in \mathcal{J}} \left(V_j(q_j) - \sum_{l \in \mathcal{L}} \lambda_l q_{j,l} \right) + \sum_{s \in \mathcal{S}} \left(\sum_{l \in \mathcal{L}} \lambda_l q_{s,l} - C_s(q_s) \right) \tag{7}$$

$$s.t. \quad q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \tag{8}$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \tag{9}$$

which is decomposable into $|\mathcal{J}| + |\mathcal{S}|$ independent sub-problems, one for each seller and buyer.

The following result states conditions under which an optimal solution of the Lagrangian dual problem and the corresponding primal solutions correspond to an efficient allocation and form a Walrasian equilibrium.

Proposition 1 *Let $\lambda^\star$ be an optimal solution of the Lagrangian dual problem, and $q^\star \in \arg\max_q \{L(q;\lambda) : q_j \in \mathcal{D}_j, j \in \mathcal{J}; q_s \in \mathcal{D}_s, s \in \mathcal{S}\}$ the optimal solution of the corresponding dual relaxation. Under conditions of:*

1. *convexity of production and consumption functions, and convexity of the feasibility sets (Assumptions (A1) and (A2)),*
2. *stability of problem (MC),*
3. *feasibility of $q^\star$ for (MC),*

$q^\star$ is a socially-efficient allocation and $\lambda^\star$ is a corresponding vector of Walrasian prices.

Proof. The social-efficiency of the allocation q^* follows immediately from the strong Lagrangian duality theorem (Geoffrion 1971), which precludes the existence of a duality gap between the primal problem and the Lagrangian dual problem under convexity and stability conditions. With respect to the fact that q^* and λ^* constitute a Walrasian equilibrium, one may simply notice that each one of the $|\mathcal{J}| + |\mathcal{S}|$ sub-problems of the Lagrangian relaxation ($\mathrm{LR}_{(\lambda)}$) corresponds to the maximization of the surplus of a seller or a buyer. ■

The *stability* of problem (MC) (Geoffrion 1971) is a key condition for the "duality gap" between the primal problem (MC) and its Lagrangian dual to vanish. Stability is garanteed, in particular, when the objective and the constraints of the primal problem are affine or quadratic. A more general condition is the *filling property* (Hiriart-Urruty and Lemaréchal 1991), which only requires compactness of the subsets $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$, and continuity of functions $V_j, j \in \mathcal{J}$ and $C_s, s \in \mathcal{S}$. This condition is sufficient to ensure the absence of duality gaps when combined with the convexity of problem (MC).

4.1 The subgradient approach

The subgradient algorithm, an approach that finds its inspiration in the gradient method for convex differentiable optimization, has been traditionally used to solve Lagrangian dual problems. At each iteration of the (basic version) of the algorithm, a subgradient of the dual function Θ is computed at the current vector of Lagrangian multipliers and the multipliers are updated along the direction of the subgradient. For the problem at hand, the subgradient algorithm can be stated as follows:

STEP 0: Set $k = 0$, $\Theta_\star^{(0)} = +\infty$, $k_\star = 0$. Initialize the vector of Lagrangian multipliers, e.g., set $\lambda^{(0)} = 0$.

STEP 1: Evaluate $\Theta(\lambda^{(k)})$. Compute a subgradient $g^{(k)} \in \partial\Theta(\lambda^{(k)})$ of Θ at $\lambda^{(k)}$:

let $q^{(k)} \in \arg\max_q \{L(q; \lambda^{(k)}) : q_j \in \mathcal{D}_j, j \in \mathcal{J}; q_s \in \mathcal{D}_s, s \in \mathcal{S}\}$;

$g^{(k)} = \sum_{s\in\mathcal{S}} q_s^{(k)} - \sum_{j\in\mathcal{J}} q_j^{(k)}$.

Update $\Theta_\star^{(k)}$: if $\Theta(\lambda^{(k)}) < \Theta_\star^{(k)}$ then set $\Theta_\star^{(k)} = \Theta(\lambda^{(k)})$, and $k_\star = k$.

STEP 2: If $g^{(k)} = 0$ then $\lambda^{(k)}$ is an optimal solution of (LD). Return $\lambda^{(k)}$ and $q^{(k)}$. Otherwise, adjust the Lagrangian multipliers according to $\lambda^{(k+1)} = \lambda^{(k)} - \xi^{(k)} g^{(k)}$.

STEP 3: If an appropriate stopping criterion is satisfied, return $\lambda^{(k)}$ and $q^{(k)}$. Otherwise, set $k = k + 1$ and return to STEP 2.

The choice of the stepsize $\xi^{(k)}$ is critical to the convergence of the algorithm. In practice, the following two schemes are the most commonly used.

1. The series $\{\xi^{(k)}\}_{k\in\mathbb{N}}$ is such that $\lim_{k\to\infty}\xi^{(k)} = 0^+$ and $\sum_{k=0}^{\infty}\xi^{(k)} = +\infty$.

2. Suppose an "estimate" $\bar{\Theta}^{(k)}$ on the optimal value of (MC) is available at each iteration k. Consider the series $\{\xi^{(k)}\}_{k\in\mathbb{N}}$ such that $\xi^{(k)} = \rho^{(k)}\frac{\Theta(\lambda^{(k)})-\bar{\Theta}^{(k)}}{\|g^{(k)}\|^2}$, where $\rho^{(k)}$ is a scaling factor. Typically, the sequence $\{\rho^{(k)}\}_{k\in\mathbb{N}}$ is such that $\rho^{(0)} \in [0,1]$ and $\rho^{(k)}$ is halved every time $\Theta_*^{(k)}$ has not been updated for n (generally equal to 10 or 20) consecutive iterations.

The determination of appropriate stopping criteria is among the weakest points of the subgradient method. A reasonable condition to verify for interrupting the algorithm would be $\|g^{(k)}\| \leq \epsilon$ (the rationale of that criterion is Everett's theorem (Everett 1963), which states that, as far as the supply/demand constraints 2 are concerned, the corresponding primal solution $q^{(k)}$ would be $\epsilon-$feasible). However, the condition $\|g^{(k)}\| \leq \epsilon$ may never be satisfied due to the fact that the subgradient algorithm only requires one subgradient of the dual function, and not the entire subdifferential, to be evaluated at each iteration. Thus, this test is often combined to other "heuristic" criteria, such as: "stop if $\Theta_*^{(k)}$ has not been improved in the last N last iterations", or "stop if the gap $\Theta(\lambda^{(k)}) - \bar{\Theta}^{(k)}$ is smaller than a threshold ϵ".

The subgradient algorithm has an interesting interpretation, illustrated in Figure 2, as an iterative auction process. The market-maker arbitrarily sets an initial vector $\lambda^{(0)}$ of single-product prices. At a given round k of the process, each seller $s, s \in \mathcal{S}$, determines a production level $q_s^{(k)}$ that maximizes its surplus given the current prices of the goods and formulates a bid $B_s^{(k)} = \{q_s^{(k)}\}$ that specifies this production level. Similarly, each buyer $j, j \in \mathcal{J}$, formulates a bid $B_j^{(k)} = \{q_j^{(k)}\}$, where $q_j^{(k)}$ is a surplus-maximizing consumption level for buyer j at the given prices. Given the bids it receives, the market-maker revises the prices of the goods along a steepest descent direction given by the excess vector $\sum_{s\in\mathcal{S}} q_s^{(k)} - \sum_{j\in\mathcal{J}} q_j^{(k)}$.

Considering that the subgradient-based auction may stop before an "implementable" outcome is reached (one that satisfies - at least approximately - the balance of supply and demand), a process for recovering feasible primal solutions is needed. In many practical cases, specialized heuristics can often be designed for that purpose. The general nature of the problem at hand, however, only allows for equally general procedures to recover primal feasibility. Hence, we adapt a very simple approach due to Larsson, Patriksson, and Strömberg (1999), which consists in *projecting* upon the feasible domain of (MC) the elements of an ergodic sequence of primal solutions converging to an optimal solution of (MC). More specifically, it can be shown that the sequence $\{\bar{q}^{(k)}\}_{k\in\mathbb{N}}$ such that $\bar{q}^{(k)} = \frac{\sum_{r=1}^{k-1}\xi^{(r)}q^{(r)}}{\sum_{r=1}^{k-1}\xi^{(r)}}$, $k \in \mathbb{N}$ converges to the set of optimal solutions of (MC). Let K be the last iteration of the auction. The (Euclidean) projection of $\bar{q}^{(K)}$ on the feasible domain of (MC) corresponds to allocation vectors q that solve:

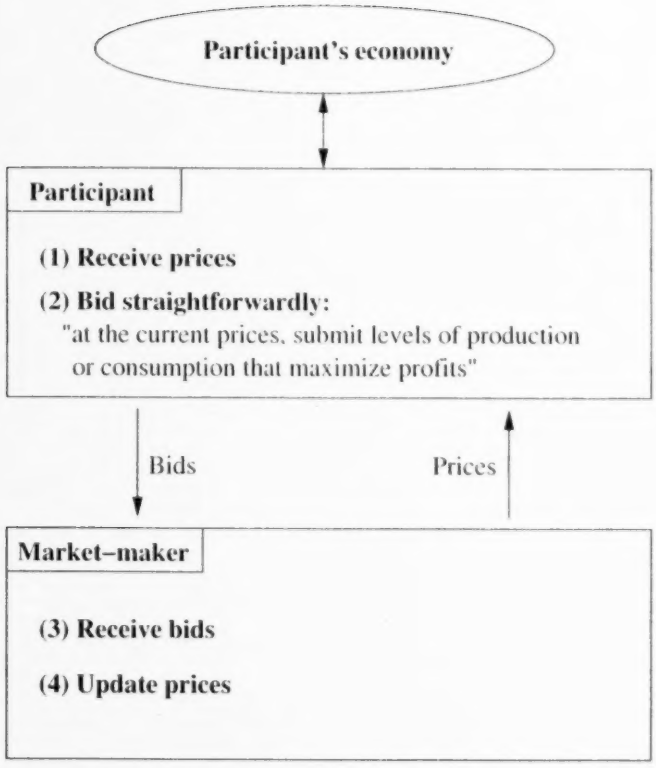

Figure 2: The subgradient algorithm as an iterative auction.

$$\min_{q} \quad \|q - \bar{q}^{(K)}\|^2 \tag{10}$$

$$s.t. \quad \sum_{j \in \mathcal{J}} q_{j,l} - \sum_{s \in \mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \tag{11}$$

$$q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \tag{12}$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \tag{13}$$

The solution to this problem is nevertheless conditionned by the availability to the market-maker of complete knowledge of feasibility sets $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$. In our case, bids submitted by sellers and buyers in previous iterations of the auction can be used to "shape" approximations of the actual feasibility sets. Consider the convex hulls:

$$\hat{\mathcal{D}}_j = \{q_j = \sum_{k=0}^{K} q_j^{(k)} \alpha_j^k : \sum_{k=0}^{K} \alpha_j^k = 1; \alpha_j^k \geq 0, \forall k = 0, \ldots, K\}, j \in \mathcal{J}$$

and

$$\hat{\mathcal{D}}_s = \{q_s = \sum_{k=0}^{K} q_s^{(k)} \beta_s^k : \sum_{k=0}^{K} \beta_s^k = 1; \beta_s^k \geq 0, \forall k = 0, \ldots, K\}, s \in \mathcal{S}$$

By virtue of the convexity of the feasibility sets, $\hat{\mathcal{D}}_j, j \in \mathcal{J}$ and $\hat{\mathcal{D}}_s, s \in \mathcal{S}$ are inner-approximations of $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$, respectively. The projection of $\bar{q}^{(K)}$ on the approximated feasible sets yields the quadratic problem:

$$\min_{q} \quad \|q - \bar{q}^{(K)}\|^2 \tag{14}$$

$$s.t. \quad \sum_{j \in \mathcal{J}} q_{j,l} - \sum_{s \in \mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \tag{15}$$

$$q_j \in \hat{\mathcal{D}}_j, \quad j \in \mathcal{J} \tag{16}$$

$$q_s \in \hat{\mathcal{D}}_s, \quad s \in \mathcal{S} \tag{17}$$

One should be warned against the "non exact" nature of this approach. Problem (14-17) is a restriction of (10-13) and, thus, is not necessarily feasible and may not be successful in providing a feasible outcome when only "few" bids are used to approximate the feasibility sets, or when the feasible domain of the allocation problem is tight.

4.2 The bundle approach

Bundle methods, originally developped for nonsmooth optimization (Wolfe 1975, Lemaréchal 1989), may equally be suggested for solving the Lagrangian dual problem. These methods

rely basically on the *bundle of information* concept, which is used to build "good" approximation models of the dual function Θ - at least in the vicinity of an optimal solution. Hence, let the bundle $\mathcal{B} = \{(\lambda^{(k)}; \Theta(\lambda^{(k)}); g^{(k)})\}_{k=1,\ldots,K}$ represent the information gathered at a given time, where $g^{(k)} \in \partial\Theta(\lambda^{(k)})$, $\forall k = 1, \ldots, K$ is a subgradient of Θ at $\lambda^{(k)}$. The first-order approximation of Θ with the information in bundle $\mathcal{B}$ yields the *cutting-plane model* of Θ: $\Theta_{\text{cp}}(\lambda) = \max_{1 \leq k \leq K}\{\Theta(\lambda^{(k)}) + g^{(k)^T}(\lambda - \lambda^{(k)})\}$.

The early cutting-plane algorithm (Kelley 1960) is an iterative procedure that consists in minimizing the approximate model Θ_{cp} and using the optimal solution $\lambda^{(K+1)}$ obtained at iteration K to enrich Θ_{cp} with a new cutting plane. Practical experience with the cutting-plane algorithm has nonetheless revealed its *instability*: the iterate $\lambda^{(K+1)}$ is often very remote from $\lambda^{(K)}$, even if the latter is very close to an optimal dual solution. A significant number of the cutting planes generated by the algorithm are consequently of little help in closing the gap between Θ and Θ_{cp} in the neighborhood of an optimal solution. Bundle methods address the instability issue by defining a *stability center* $\bar{\lambda}$ and requiring that the approximate model produces an iterate $\lambda^{(K+1)}$ "not too far" from $\bar{\lambda}$. This is done by the introduction of a stabilizing term $\frac{1}{2t^K}\|\lambda - \bar{\lambda}\|^2$ in the expression of the cutting-plane model Θ_{cp}, where t^K is a parameter that can be interpreted both as a stepsize and a *trust-region* parameter. The new approximation model of Θ is thus $\Theta_B(\lambda) = \Theta_{\text{cp}}(\lambda) + \frac{1}{2t^K}\|\lambda - \bar{\lambda}\|^2$. The minimization of Θ_B at iteration K corresponds to the quadratic problem $(Q_{\mathcal{B}})$:

$$\min_{\nu,\lambda} \quad \nu + \frac{1}{2t^K}\|\lambda - \bar{\lambda}\|^2 \tag{18}$$

$$s.t. \quad \nu \geq \Theta(\lambda^{(k)}) + g^{(k)^T}(\lambda - \lambda^{(k)}), \quad k = 1, \ldots, K \tag{19}$$

Let $(\nu^{(K+1)}; \lambda^{(K+1)})$ be the optimal solution of this problem, and let $\Delta^K = \Theta(\bar{\lambda}) - \Theta(\lambda^{(K+1)})$ and $\tilde{\Delta}^K = \Theta(\bar{\lambda}) - \Theta(\lambda^{(K+1)})$ denote the *actual* and the *predicted* (by model Θ_B) decrease of Θ, respectively. If $\Delta^K \geq m\tilde{\Delta}^K$ (m is a pre-specified parameter such that $0 < m < 1$), i.e., the value of Θ has actually been "sufficiently" decreased with respect to the predicted value, the bundle method performs a *serious-step*: accept $\lambda^{(K+1)}$ as the new stability center. Otherwise, a *null-step*, which consists in leaving the stability center unchanged but adding $(\lambda^{(K+1)}; \Theta(\lambda^{(K+1)}); g^{(K+1)})$ to the bundle for a more refined approximation of Θ, is made.

While it does not affect the validity of the approach, the setting of the parameter sequence $\{t_k\}_{k\in\mathbb{N}}$ plays a decisive role in the behavior of a bundle algorithm and its numerical efficiency. Theorical and practical evidence shows the difficulties that arise from setting a *fixed* parameter $t_k = t, \forall k \in \mathbb{N}$ (Hiriart-Urruty and Lemaréchal 1991). Hence, a small value of t will drive the bundle algorithm into making relatively few null-steps and also "small" serious-steps resulting only in marginal improvement of the dual functions (when $t \to 0$, the bundle method is nothing else than the subgradient algorithm). On the other hand, the bundle algorithm performs few serious-steps when large values of t are considered, moving toward the cutting-plane algorithm as $t \to \infty$. The design of variable sequences $\{t_k\}_{k\in\mathbb{N}}$ is indeed

a complex issue and the literature is clearly lacking in theorical results on provably "good" sequences. To date, heuristic approaches that consist in increasing t_k after a serious-step and decreasing it after a null-step (Kiwiel 1990, Schramm and Zowe 1992) seem to provide the best results.

It is interesting to compare the way bundle methods manage prices with the simpler price update scheme of the subgradient algorithm. Basically, two fundamental observations can be made: (a) the bundle approach relies on a collection of information representing a "history" of the market, that is, a set of prices and the corresponding bidder reactions (the desired production and consumption levels at these prices); (b) a specific price vector (the stability center), in the neighborhood of which the approximate cutting-plane model can be reasonably "trusted", is given a special status. In that regard, the dual viewpoint provides additional insight. The Kuhn-Tucker optimality conditions for $Q_{\mathcal{B}}$ imply that, for an optimal solution $(\nu^\star; \lambda^\star)$ of $Q_{\mathcal{B}}$, there exists a vector of multipliers $\{\delta_k^\star\}_{k=1,\ldots,K}$ such that:

(i) $d^\star = \lambda^\star - \bar{\lambda} = -t^K \sum_{k=1}^K \delta_k^\star g^{(k)}$;

(ii) $\nu^\star = -t^K \| \sum_{k=1}^K \delta_k^\star g^{(k)} \|^2 - \sum_{k=1}^K \delta_k^\star c_k + \Theta(\bar{\lambda})$, where $c_k = \Theta(\bar{\lambda}) - (\Theta(\lambda^{(k)}) + g^{(k)^T}(\bar{\lambda} - \lambda^{(k)}))$, $k = 1, \ldots, K$;

(iii) $\delta_k^\star(-\nu^\star + g^{(k)^T}(\lambda^\star - \lambda^{(k)}) + \Theta(\lambda^{(k)})) = 0$, $k = 1, \ldots, K$;

(iv) $\sum_{k=1}^K \delta_k^\star = 1$ and $\delta_k^\star \geq 0, k = 1, \ldots, K$.

Conditions (i) and (iv) are particularly instructive. Together, they indicate that bundle methods actually construct *aggregated subgradients* $\tilde{g}^{(K)} = \sum_{k=1}^K \delta_k^\star g^{(k)}$ as convex combinations of the subgradients available in the bundle, and move (in the case of a serious-step) in the opposite direction of the aggregated subgradient, to an extent given by stepsize t^K.

5 An auction scheme based on column-generation

Another decomposition technique that can be given a market interpretation as an iterative auction is the Dantzig-Wolfe decomposition principle. Let us suppose sets of feasible consumption levels $\{q_j^r\}_{r \in k_j}$, $j \in \mathcal{J}$, and production levels $\{q_s^r\}_{r \in k_s}$, $s \in \mathcal{S}$ are available to the market-maker. Consider the corresponding convex hulls:

$$\tilde{\mathcal{D}}_j^{k_j} = \{q_j = \sum_{r=1}^{k_j} q_j^r \alpha_j^r : \sum_{r=1}^{k_j} \alpha_j^r = 1; \alpha_j^r \geq 0, \forall r = 1, \ldots, k_j\}, j \in \mathcal{J}$$

and

$$\tilde{\mathcal{D}}_s^{k_s} = \{q_s = \sum_{r=1}^{k_s} q_s^r \beta_s^r : \sum_{r=1}^{k_s} \beta_s^r = 1; \beta_s^r \geq 0, \forall r = 1, \ldots, k_s\}, s \in \mathcal{S}$$

Inner-linearization (Geoffrion 1970) of the production (consumption) feasibility sets and the cost (valuation) functions suggests to solve linear programming approximations of the non-linear market-clearing problem (MC). More precisely, it proceeds in the following two steps:

- Approximate $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$ with $\tilde{\mathcal{D}}_j^{k_j}$ and $\tilde{\mathcal{D}}_s^{k_s}$, respectively.
- For $q_j \in \mathcal{D}_j, j \in \mathcal{J}$, replace $V_j(q_j) = V_j(\sum_{r=1}^{k_j} q_j^r \alpha_j^r)$ with $\sum_{r=1}^{k_j} \alpha_j^r V_j(q_j^r)$. Similarly, replace $C_s(q_s) = C_s(\sum_{r=1}^{k_s} q_s^r \beta_s^r)$ with $\sum_{r=1}^{k_s} \beta_s^r C_s(q_s^r)$.

These approximations yield the restricted master problem (MC-R):

$$\max \quad \sum_{j\in\mathcal{J}} \sum_{r=1}^{k_j} \alpha_j^r V_j(q_j^r) - \sum_{s\in\mathcal{S}} \sum_{r=1}^{k_s} \beta_s^r C_s(q_s^r) \tag{20}$$

$$s.t. \quad \sum_{j\in\mathcal{J}} \sum_{r=1}^{k_j} \alpha_j^r q_{j,l}^r - \sum_{s\in\mathcal{S}} \sum_{r=1}^{k_s} \beta_s^r q_{s,l}^r = 0, \quad l \in \mathcal{L} \tag{21}$$

$$\sum_{r=1}^{k_j} \alpha_j^r = 1, \quad j \in \mathcal{J} \tag{22}$$

$$\sum_{r=1}^{k_s} \beta_s^r = 1, \quad s \in \mathcal{S} \tag{23}$$

$$\alpha_j^r \geq 0, \quad r = 1, \ldots k_j, \quad j \in \mathcal{J} \tag{24}$$

$$\beta_s^r \geq 0, \quad r = 1, \ldots k_s, \quad s \in \mathcal{S} \tag{25}$$

Let $[\{\tilde{\alpha}_j^r\}_{r=1,\ldots,k_j, j\in\mathcal{J}}; \{\tilde{\beta}_s^r\}_{r=1,\ldots,k_s, s\in\mathcal{S}}]$ be an optimal basic solution of (MC-R) and $\{\tilde{\mu}_l\}_{l\in\mathcal{L}}$, $\{\tilde{\tau}_j\}_{j\in\mathcal{J}}$, and $\{\tilde{\tau}_s\}_{s\in\mathcal{S}}$ the corresponding simplex multipliers. The relative cost optimality conditions for the restricted master problem are:

$$V_j(q_j^r) - \sum_{l\in\mathcal{L}} \tilde{\mu}_l q_{j,l}^r - \tilde{\tau}_j \leq 0, \quad \forall r = 1, \ldots, k_j, \forall j \in \mathcal{J}$$

and

$$-C_s(q_s^r) + \sum_{l\in\mathcal{L}} \tilde{\mu}_l q_{s,l}^r - \tilde{\tau}_s \leq 0, \quad \forall r = 1, \ldots, k_s, \forall s \in \mathcal{S}$$

Thus, the generation of new feasible consumption levels $q_j^{k_j+1}$, $j \in \mathcal{J}$, or production levels $q_s^{k_s+1}$, $s \in \mathcal{S}$ that eventually improve the approximation can be done by pricing out sets $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$; that is, by solving the subproblems

$$(\mathrm{SP})_j: \quad \max_{q_j \in D_j} \{V_j(q_j) - \sum_{l \in \mathcal{L}} \tilde{\mu}_l q_{jl} - \tilde{\tau}_j\}, j \in \mathcal{J};$$

and

$$(\mathrm{SP})_s: \quad \max_{q_s \in D_s} \{-C_s(q_s) + \sum_{l \in \mathcal{L}} \tilde{\mu}_l q_{sl} - \tilde{\tau}_s\}, s \in \mathcal{S}.$$

The Dantzig-Wolfe decomposition principle has a classical economic interpretation as a decentralized planning procedure (Dantzig 1963, chapter 23). A central authority (the headquarters) has to devise an optimal operation plan for an enterprise composed of several subsidiaries. Each subsidiary has private information concerning its technology and how it constrains its contribution to the overall plan. The headquarters deals with the constraints concerning the resource exchange between the subsidiaries. The Dantzig-Wolfe decomposition principle can be viewed as an iterative decision process in which the role of the central authority is to determine an optimal operation plan given a set of partial operation plans suggested by the subsidiaries and to announce corresponding simplex multipliers (interpreted as prices), while the subsidiaries react to the announced prices by proposing new promising partial plans.

In the specific context of the exchange economy considered in this paper, it turns out that this principle translates easily into a two-phase auction:

Phase 1.

In this initialization phase, the market-maker asks the sellers and the buyers to submit the $|\mathcal{L}| + |\mathcal{J}| + |\mathcal{S}|$ bids required to build a first restricted master problem.

Denote by $k_j^{(0)}$ the number of bids submitted by buyer j during the initialization phase, and by $k_s^{(0)}$ the number of bids submitted by seller s.

Set $n = 0$.

Phase 2.

At iteration n:

- Suppose $k_j^{(n)}$ and $k_s^{(n)}$ bids have been submitted, up to iteration n, by buyer j and seller s, respectively. The market-maker solves the restricted master problem (MC-R) and announces prices $\{\tilde{\mu}_l\}_{l \in \mathcal{L}}$, as well as multipliers $\{\tilde{\tau}_j\}_{j \in \mathcal{J}}$ and $\{\tilde{\tau}_s\}_{s \in \mathcal{S}}$ to the participants.
- Each buyer $j, j \in \mathcal{J}$, determines a surplus-maximizing solution $q_j^{k_j^{(n)}}$ of $(\mathrm{SP})_j$. If $q_j^{k_j^{(n)}}$ improves on the approximation, buyer j submits bid $B_j^{(n)} = \{q_j^{k_j^{(n)}}; V(q_j^{k_j^{(n)}})\}$.

- Similarly, each seller s, $s \in \mathcal{S}$, determines a solution $q_s^{k^{(n)}}$ of $(\mathrm{SP})_s$. If $q_s^{k^{(n)}}$ improves the approximation, seller s submits bid $B_s^{(n)} = \{q_s^{k^{(n)}} : C(q_s^{k^{(n)}})\}$.

Although Lagrangian relaxation techniques and the Dantzig-Wolfe decomposition method can be seen as duals of each other (see, for example, Lasdon 1970, chapter 8), there are important differences between the corresponding auction schemes. First, from an *informational* point of view, relaxation-based auctions are *punctual* mechanisms regarding the determination of allocations, in the sense that only the last bids made by the participants are relevant to the market-maker. By opposition, the Dantzig-Wolfe auction aggregates all the bids it receives into a feasible allocation and may thus be considered as *cumulative* process. A straightforward consequence of the latter point is that the prices announced by the market-maker often do not form an equilibrium with any allocation corresponding to optimal solutions of sellers' and buyers' surplus-maximization subproblems. As previously noticed, bundle methods use more sophisticated pricing schemes in comparison with the subgradient algorithm, exploiting a history of the market made of all the prices computed during a certain time window and the corresponding bidder reactions to these prices. It is also noteworthy that the Dantzig-Wolfe auction process, being a restriction-based approach, is guaranteed to produce feasible allocations at any time, which is not the case of the Lagrangian methods that require an extra treatment of the allocations obtained (e.g., projection on the domain of feasible allocations or an approximation of the latter, see section 4) to be done.

The *strategic* viewpoint is equally interesting. One could argue that auction schemes based on Lagrangian relaxation are iterative mechanisms that determine socially efficient allocations under the assumption that the sellers and buyers follow *myopic* best-response bidding strategies. That is, they bid for production and consumption levels that maximize their surplus given the current prices announced by the market-maker. Thus, a participant does not consider future rounds of the auction, nor does it exploits "what it knows" about other participants' private data. The Dantzig-Wolfe auction process, on the other hand, requires the participants to formulate bids that contain, in addition to the desired production and consumption levels at the prices, cost and valuation information. The fact that the participants may misrepresent their true costs and valuations, while continuing to bid on surplus-maximizing quantities, implies that even the myopic best-response strategy is too weak to ensure that a Dantzig-Wolfe auction will produce socially efficient final outcomes.

6 Computational study

6.1 The experimental setting

In order to illustrate the above-mentioned auction mechanisms, we call upon a more detailed model of multi-lateral multi-commodity markets presented in Bourbeau *et al.* (2004), which has the advantage of being closer to actual applications, especially in the context of regulated

marketplaces for the trade of natural resources. We briefly present in the following the notation and the important elements of the model.

Participants in the market seek to trade a set of *products*. A product is a basic commodity with a specific physical denomination (e.g., a wood specie). Products are not available in a "pure" state, however, but come rather as part of different *lots* that are "mixtures" of several products. Hence, let

- $\mathcal{K}$: set of basic products.
- $\mathcal{L}$: set of lots.
- b_l^k: the proportion of product k in lot l, $k \in \mathcal{K}$, $l \in \mathcal{L}$.

It is assumed for simplicity that each seller may only procure a single lot. Thus, a lot $l \in \mathcal{L}$ is attached to seller l and Q^l denotes the maximum quantity produced of that lot. The production cost function $C_l(.)$ of lot l is assumed to be such that its corresponding marginal cost function $C_l'(.)$ is continuous, piecewise-linear, and strictly increasing (Bourbeau *et al.* 2004).

On the buyer side, Bourbeau *et al.*'s model accounts for the differences in quality among the various lots by considering: (i) a multiplicative adjustment coefficient r_j^l, which indicates that one unit of lot l is equivalent for buyer j to r_j^l units of a standard lot; and (ii) an additive coefficient s_j^l, which denotes how much more or less buyer j valuates, in absolute terms, a unit of of lot l with respect to a unit of the standard lot. Furthermore, the model considers a unit transportation cost t_j^l between the seller producing lot l and buyer j. The latter's preference for a bundle $q_j = \{q_{j,l}\}_{l \in \mathcal{L}}$ can accordingly be expressed as $V_j(q_j) = U_j(\sum_{l \in \mathcal{L}} r_j^l q_{j,l}) + \sum_{l \in \mathcal{L}} (s_j^l - t_j^l) q_{j,l}$, where $U_j(.)$ is a utility function such that $U_j'(.)$ is continous, piecewise-linear, and strictly decreasing. Buyers need also to express requirements regarding the composition of the lots they purchase. More specifically, let M_j^k and m_j^k denote respectively the maximum and minimum proportions of product k that buyer j may tolerate in the acquired lots, and Q^j the maximum total volume - expressed in terms of the standard lot - buyer j requires.

With the notation above, the market-clearing problem corresponds to the following formulation:

$$\max \quad \sum_{j\in\mathcal{J}} U_j(\sum_{l\in\mathcal{L}} r_j^l q_{j,l}) + \sum_{l\in\mathcal{L}} (s_j^l - t_j^l) q_{j,l} - \sum_{l\in\mathcal{L}} C_l(q_l) \tag{26}$$

$$s.t. \quad \sum_{j\in\mathcal{J}} q_{j,l} - q_l = 0, \quad l \in \mathcal{L} \tag{27}$$

$$q_l \leq Q^l, \quad l \in \mathcal{L} \tag{28}$$

$$\sum_{l\in\mathcal{L}} r_j^l q_{j,l} \leq Q^j, \quad j \in \mathcal{J} \tag{29}$$

$$m_j^k \sum_{l\in\mathcal{L}} r_j^l q_{j,l} \leq \sum_{l\in\mathcal{L}} b_l^k r_j^l q_{j,l} \leq M_j^k \sum_{l\in\mathcal{L}} r_j^l q_{j,l}, \quad j \in \mathcal{J}, k \in \mathcal{K} \tag{30}$$

$$q_{j,l} \geq 0, \quad j \in \mathcal{J}, l \in \mathcal{L} \tag{31}$$

$$q_l \geq 0, \quad l \in \mathcal{L} \tag{32}$$

where $q_{j,l}$ denotes the quantity of lot l purchased by buyer j and q_l the total quantity of lot l procured by the corresponding seller.

The computational experiments we realized were guided by the general objective of comparing auction processes based on various decomposition methods from the perspective of the economic efficiency achieved. The "benchmark" we used was a mechanism based on the centralized market-clearing formulation (26-32), which assumes that the market-maker has access to complete information about the sellers and buyers valuation and cost functions, as well as their technological constraints. The experiments were carried out on a 64-processor, 64 Gigabytes of RAM Sun Enterprise 10000 operated under SunOS 5.8, with versions 8.0 and 1.2 of the CPLEX solver and the Concert library, respectively.

We have performed tests on several problem series made of instances obtained from a custom problem generator we have developed. Given values for the numbers of buyers, sellers (lots), and basic products, volumes $\{Q_j\}_{j\in\mathcal{J}}$ and $\{Q_l\}_{l\in\mathcal{L}}$, proportions $\{b_l^k\}_{l\in\mathcal{L},k\in\mathcal{K}}$, and tolerances $M_j^k, m_j^k, j \in \mathcal{J}, k \in \mathcal{K}$ are randomly generated according to uniform distributions over prespecified intervals. For the sake of simplicity, we considered *purely quadratic* buyer utility functions $U_j(.), j \in \mathcal{J}$ and seller cost functions $C_l(.), l \in \mathcal{L}$. This implies no loss of generality, since a simple transformation suggested in Bourbeau *et al.* (2004) makes it possible to deal with a general piecewise-quadratic formulation as a purely quadratic one. Furthermore, our instances involved no transportation costs t_l^j or additive adjustment coefficients s_l^j, $j \in \mathcal{J}, l \in \mathcal{L}$. Table 1 displays the characteristics of the problem series considered in the study. The series are subdivided into two categories according to (1) $|\mathcal{K}|$, the number of basic products; and (2) Δ_m, the minimum difference between tolerances M_j^k, m_j^k ($\Delta_m = \underline{M} - \overline{m}$, where $\underline{M}$ designates the minimum value M_j^k can take, and $\overline{m}$ the maximum value of m_j^k). These two parameters are important since they directly impact the number and forcefulness of constraints (30) in the market-clearing formulation. Each series of problems consisted of 10 randomly generated instances.

We have set up four auction processes. The first three are based on variants of the subgradient

Problem series	Problem description			
	# buyers	# lots	# products	$\Delta_m = \underline{M} - \overline{m}$ (%)
$S-01$	50	100	3	30
$S-02$	50	250	3	30
$S-03$	100	50	3	30
$S-04$	100	200	3	30
$S-05$	100	300	3	30
$S-06$	50	100	10	10
$S-07$	50	250	10	10
$S-08$	100	50	10	10
$S-09$	100	200	10	10
$S-10$	100	300	10	10

Table 1: Characteristics of problem instances

algorithm, while the fourth relies on Dantzig-Wolfe decomposition. More specifically, the subgradient variants used are:

1. The "basic" subgradient method: $\lambda^{(k+1)} = \lambda^{(k)} - \xi^{(k)} g^{(k)}$, with stepsize formula $\xi^{(k)} = \rho^{(k)} \frac{\Theta(\lambda^{(k)}) - \bar{\Theta}^{(k)}}{\|g^{(k)}\|^2}$. We have used the simple estimate $\bar{\Theta}^{(k)} = 0.5.\Theta_*^{(k)}$, where $\Theta_*^{(k)}$ denotes the best value of the lagrangian dual function $\Theta(.)$ found so far. Parameter $\rho^{(0)}$ has been calibrated in the set $\{0.1, 0.3, 0.5, 0.7, 1.0\}$ for each problem series.

2. The subgradient method with the Camerini-Fratta-Maffioli rule (Camerini, Fratta, and Maffioli 1975). This variant relies on an elementary aggregation of the subgradients to compute a direction along which to move at each iteration. Thus, $\lambda^{(k+1)} = \lambda^{(k)} - \xi^{(k)} d^{(k)}$, where $d^{(k)} = g^{(k)} + \sigma^{(k)} d^{(k-1)}$ and $\sigma^{(k)}$ is such that

$$\sigma^{(k)} = \begin{cases} -\mu \frac{g^{(k)} d^{(k-1)}}{\|d^{(k)}\|^2} & \text{if} g^{(k)} d^{(k-1)} < 0, \\ 0 & \text{otherwise,} \end{cases}$$

   where parameter μ is set to 1.5. The stepsize formula used is similar to that of the basic subgradient, that is $\xi^{(k)} = \rho^{(k)} \frac{\Theta(\lambda^{(k)}) - \bar{\Theta}^{(k)}}{g^{(k)} d^{(k)}}$.

3. The subgradient method with the modified Camerini-Fratta-Maffioli rule. Instead of using a "fixed" value of parameter μ, this variant uses the rule: $\mu = -\frac{\|g^{(k)}\| \|d^{(k-1)}\|}{g^{(k)} d^{(k-1)}}$.

In order to compare the different auction mechanisms, we fixed the maximum number of rounds to 1000 for the subgradient-based auctions, and to 200 for the DW auction. The following metrics are used:

(a) For the auction processes based on the subgradient and its variants (CFM and modified CFM), we measured:

1. the gap $GAP^{\prime}_{\rm lr} = (Z_{\rm lr} - Z_{\rm cent})/Z_{\rm cent}$, where $Z_{\rm lr}$ is the best upper bound obtained by the corresponding subgradient methods;
2. the gap $GAP^{P^{\prime}}_{\rm lr} = (Z_{\rm cent} - \hat{Z})/Z_{\rm cent}$, where $\hat{Z}$ is the economic surplus achieved by the "closest" feasible allocation to the primal solution $q^{(k_\bullet)}$, obtained by projecting the latter on the feasible domain of model (26-32);
3. the gap $GAP^{E}_{\rm lr}$ corresponding to the allocation q^E obtained through the projection of the last term $q^{(K)}$ of the ergodic sequence $\{\bar{q}^{(k)}\}_{k\in\mathbb{N}}$ defined in Section 4.1 on the domain of feasible allocations, that is $GAP^{E}_{\rm lr} = (Z_{\rm cent} - \bar{Z})/Z_{\rm cent}$, where $\bar{Z}$ is the economic surplus achieved by q^E;
4. the Euclidean norms $\|g^{(k_\bullet)}\|$ and $\|g^{(K)}\|$ of the trivial sub-gradients corresponding to allocations $q^{(k_\bullet)}$ and $q^{(K)}$, respectively.

(b) For the DW-based auction process, we evaluated the gap $GAP_{\rm dw} = (Z_{\rm cent} - Z_{\rm dw})/Z_{\rm cent}$ between the optimal value $Z_{\rm cent}$ obtained by the centralized allocation model (26-32) and the lower bound obtained by Dantzig-Wolfe auction process $Z_{\rm dw}$.

6.2 Numerical results

Table 2 displays the results obtained by the basic subgradient, the Camerini-Fratta-Maffioli, and the modified Camerini-Fratta-Maffioli methods. The first figure ($GAP^{\prime}_{\rm lr}$) indicates the average gaps corresponding to the best upper bound achieved at the last (1000th) bidding round. We only retained the best gaps (with respect to the five possible values of parameter $\rho^{(0)}$) and we listed the best $\rho^{(0)}$s in the table. The second and the third entries of the table are the average gaps of the allocations obtained by projecting $q^{(k_\bullet)}$ and $q^{(K)}$, respectively, on the feasible allocation domain. The fourth and fifth figures indicate the average norms of the subgradients corresponding to $q^{(k_\bullet)}$ and $q^{(K)}$.

Several interesting observations can be made regarding these results. First, while both the basic subgradient and the CFM methods display quite small average gaps, consistently converging to within 3% of the optimal solution of the centralized market-clearing formulation, the convergence of the modified CFM is a more mixed bag: on some series ($S-01$ and $S-06$, for instance), the average gaps are comparable to those of the two other methods; on others (especially $S-03$, $S-07$, and $S-08$) the gaps are much larger. In that regard, the poor performance of the modified CFM method seems likely to be a consequence of its relatively greater sensitivity to the choice of the initial scaling factor $\rho(0)$, rather than an inherent lack of effectiveness. The feasibility of the primal solutions returned by the three methods is a major source of concern, though, as fairly large subgradients (with respect of the magnitude of the randomly generated quantities $Q^j, j \in \mathcal{J}$ and $Q^l, l \in \mathcal{L}$) are displayed, and the gaps associated with the projection of the primal solution on the feasible allocation domain are large too. It is noticeable that, despite significantly lowering infeasibility, the ergodic sequence of Larsson, Patriksson, and Strömberg (1999) does not solve this issue in a satisfactory manner, as their primal convergence seems to be extremely slow, while the

projection of the terms of the sequence on the feasible domain yields even larger gaps than the original allocations.

In order to gain more insight into the behavior of the three relaxation-based auctions, we have taken the third instance of series $S-01$ and we mapped out in Figure 3 the best upper bound of each auction as the latter evolves in time. The figure shows quite clearly the relatively superior speed of convergence of the two CFM methods, which was expected: whereas barely 189 and 183 rounds of bidding are enough for the CFM and the modified CFM methods, respectively, to attain a less than 1% gap, the basic subgradient method needs 463 rounds to achieve that same level. Another issue which is interesting to examine in detail is the sensitivity of the subgradient methods to the initial scaling factor $\rho^{(0)}$. We have thus considered another problem in series $S-01$ and plotted in Figure 4 the upper bound of the three methods with respect to different values of parameter $\rho^{(0)}$. The resulting plots confirm the greater sensitivity to $\rho^{(0)}$ of the modified CFM, which might perform very poorly and fail to converge within an acceptable number of rounds if the "wrong" value of the initial scaling factor is chosen (e.g., $\rho^{(0)} = 1.0$ for the instance of Figure 4).

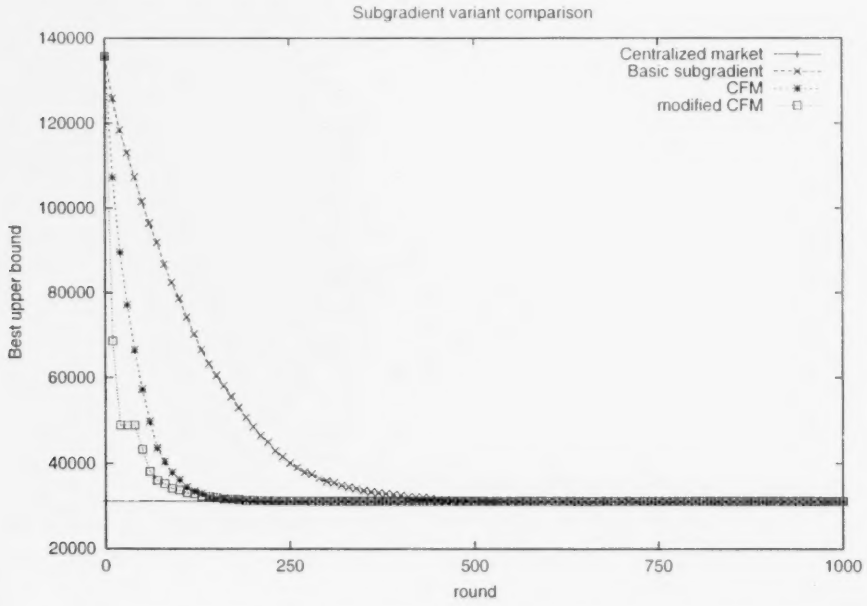

Figure 3: Evolution of the best upper bound for the three subgradient methods.

Figure 5 presents the evolution of GAP_{dw}, the gap between the lower bound at each round of the DW-based auction and the optimal surplus of the centralized market-clearing model, for instances of series $S-01$, $S-04$, and $S-09$. The initialization phase consisted in announcing a 0 price vector, then a *"large"* price vector, and collecting the bids made by the sellers and the buyers at these prices. The figure indicates that the DW process has each time succeeded, in a fairly small number of rounds, in constructing very close approximations of the actual valuation and cost functions of the participants. One could also remark that the smaller the instance ($S-01$ vs $S-04$) and the fewer products it involves ($S-04$ vs

Series	Basic subgradient					
	GAP_{lr} (%)	GAP^{P}_{lr} (%)	GAP^{E}_{lr} (%)	$\|g^{(k_\bullet)}\|$	$\|g^{(K)}\|$	$\rho_\bullet^{(0)}$
$S-01$	0.11	2.84	17.68	954.51	260.38	0.5
$S-02$	0.69	5.15	57.41	1480.18	243.71	0.5
$S-03$	1.29	12.07	76.80	1663.27	185.89	0.5
$S-04$	0.05	3.17	27.09	2118.53	442.49	0.7
$S-05$	1.76	10.47	68.44	3247.60	378.24	0.5
$S-06$	0.23	3.78	24.85	1337.48	328.63	0.7
$S-07$	2.99	15.37	70.25	2131.56	276.35	0.5
$S-08$	1.70	14.69	78.53	2342.54	190.53	0.7
$S-09$	1.21	6.28	43.84	3071.38	578.84	0.7
$S-10$	4.05	20.22	73.10	4323.10	401.54	0.7
	CFM					
	GAP_{lr} (%)	GAP^{P}_{lr} (%)	GAP^{E}_{lr} (%)	$\|g^{(k_\bullet)}\|$	$\|g^{(K)}\|$	$\rho_\bullet^{(0)}$
$S-01$	0.10	2.58	24.81	925.06	311.12	0.3
$S-02$	0.71	5.46	58.97	1509.12	247.88	0.5
$S-03$	1.22	11.32	76.56	1691.69	185.73	0.5
$S-04$	0.12	3.14	27.07	1999.62	441.17	0.7
$S-05$	1.57	9.29	67.85	3115.90	377.95	0.5
$S-06$	0.39	3.93	26.61	1302.91	342.53	0.7
$S-07$	2.49	13.40	70.05	2055.00	275.99	0.5
$S-08$	1.88	16.43	79.72	2347.58	192.34	0.7
$S-09$	1.54	7.08	44.06	3055.72	580.74	0.7
$S-10$	4.03	20.32	74.18	4314.83	404.95	0.7
	Modified CFM					
	GAP_{lr} (%)	GAP^{P}_{lr} (%)	GAP^{E}_{lr} (%)	$\|g^{(k_\bullet)}\|$	$\|g^{(K)}\|$	$\rho_\bullet^{(0)}$
$S-01$	0.08	2.40	39.92	910.22	412.35	0.1
$S-02$	0.78	5.35	62.19	1552.66	262.85	0.3
$S-03$	5.61	35.75	85.58	1788.30	199.65	0.3
$S-04$	0.16	3.22	39.15	2036.43	545.61	0.3
$S-05$	2.11	10.33	74.04	3234.97	409.98	0.3
$S-06$	0.26	3.55	39.14	1287.81	430.49	0.3
$S-07$	11.27	41.38	84.87	2318.16	311.49	0.3
$S-08$	23.30	76.02	91.69	2566.29	211.83	0.5
$S-09$	3.18	8.84	62.98	3209.39	720.68	0.3
$S-10$	28.32	67.58	92.59	4896.05	455.18	0.3

Table 2: Behavior of the subgradient, the CFM, and the modified CFM methods

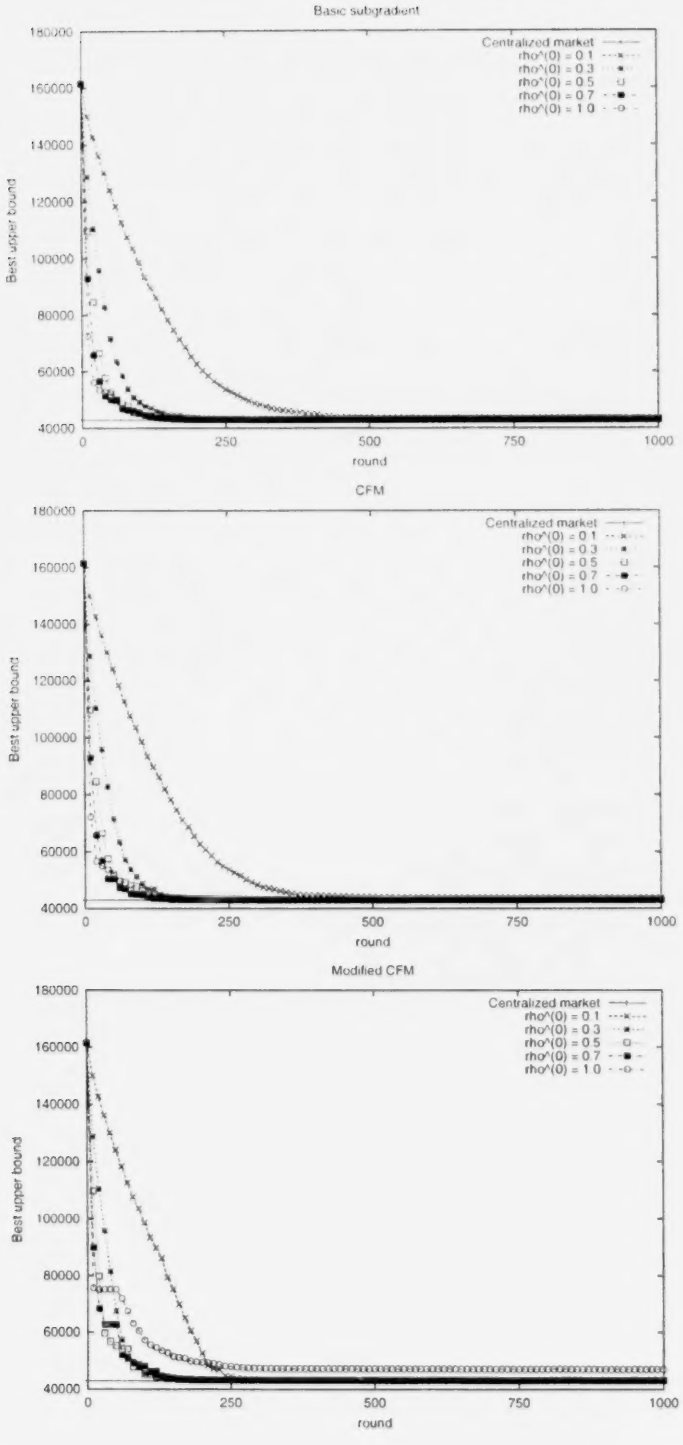

Figure 4: Sensitivity of the subgradient methods to $\rho^{(0)}$.

$S-09$), the quicker the convergence of the iterative process to the centralized market-clearing outcome.

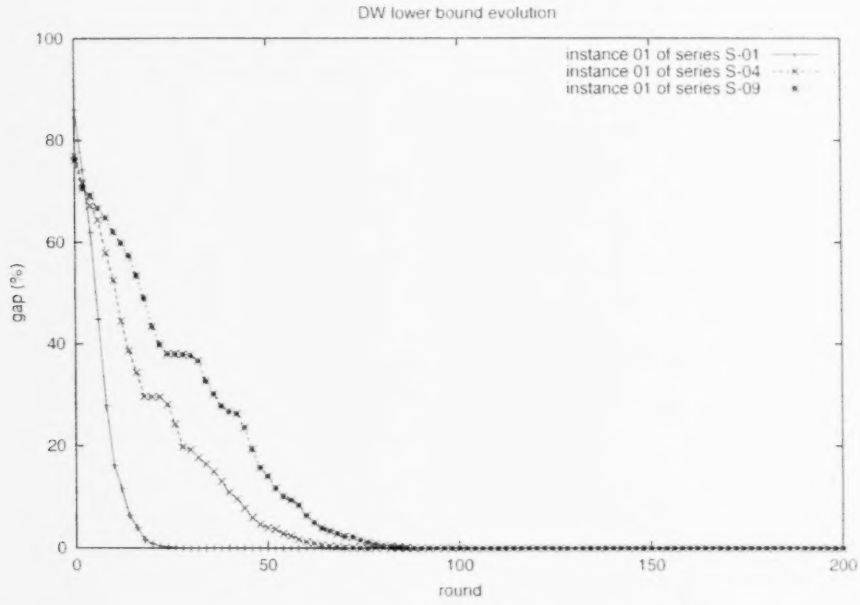

Figure 5: Evolution of the gap associated with the DW-based auction for instances of $S-01$, $S-04$, and $S-09$.

7 Conclusions

This paper has presented a new perspective of mathematical decomposition methods as iterative auctions in combinatorial exchanges of interdependent goods. We focussed on two well-known approaches, the Lagrangian relaxation and Dantzig-Wolfe decomposition, and we have shown that they can be interpreted as auction processes in which the participants progressively reveal their preferences to the market-maker. Under certain conditions, these auctions yield outcomes that reconcile the overall welfare-maximization market objective with the individual views of participants seeking to maximize their surplus. We have furthermore established that subtle yet important differences exist between the two kinds of approaches in the structure of information that the participants in the market need to disclose and how that information is exploited, as well as in the assumption made regarding the strategic behavior of the participants.

The numerical results have shown that the different variants of the subgradient method often converge in the dual space to the optimal market surplus. Nevertheless, they generally fail to produce feasible allocations. This is a major impediment to the acceptance and usability of the auction processes associated with these methods. On the other hand, Dantzig-Wolfe

decomposition maintains primal feasibility but requires the participants reveal parts of their private preferences, which may not be acceptable to them. The bundle methods, with their more sophisticated price update rules, could help get the best of both worlds, that is, to resolve the primal infeasibility issue while having minimal preference revelation requirements.

Several interesting research issues could be addressed in the future. First, the convexity assumption we made about feasibility sets of participants is a key one. In particular, when indivisible goods are considered, duality gaps prevent the interpretation of dual multipliers as prices. Two avenues that seem to be attractive are: (a) the exploration of extended formulations of the market-clearing allocation problem (see Bikhchandani and Ostroy [9] for the CAP); and (b) pricing schemes based on approximated linear prices, which sacrify either dual feasibility or complementary slackness (e.g. DeMartini *et al.* [16]). Second, we have only considered a minimal set of constraints on the market side (demand and supply balance). Real world markets would typically add other constraints derived from specific business rules, such as buyers requiring to be matched with a few "qualified" sellers, and the decomposition approaches need to be adequately adapted to deal with the additional constraints. Finally, incentive compatibility of the auction mechanisms associated with the decomposition approaches is an important and challenging issue we plan to explore.

Acknowledgments

We would like to thank Benoît Bourbeau who helped in coding the numerical methods and in realizing the experimental study. This project has received financial support from the Bell Canada University Laboratories (LUB), the Natural Sciences and Engineering Research Council of Canada through their Basic Research and Collaborative Industry-University Research grant programs, and the FCAR research development fund of the Province of Québec.

References

[1] Arrow, K. and Debreu, G. Existence of Equilibrium for a Competitive Economy. *Econometrica*, 22:265–290, 1954.

[2] Arrow, K. and Hahn, F.H. *General Competitive Analysis.* Holden-Day, San Francisco, 1971.

[3] Ausubel, L.M. An Efficient Dynamic Auction for Heterogeneous Commodities. Technical report, Department of Economics, University of Maryland, 2002.

[4] Bertsekas, D. A New Algorithm for the Assignment Problem. *Mathematical Programming*, 21:152–171, 1981.

[5] Bertsekas, D. *Linear Network Optimization.* The MIT Press, Cambridge, 1991.

[6] Bertsekas, D. Auctions Algorithms for Network Flow Problems: A Tutorial Introduction. *Computational Optimization and Applications*, 1:7–66, 1992.

[7] Bikhchandani, S., De Vries, S., Schummer, J., and Vohra, R.V. Linear Programming and Vickrey Auctions. In *Mathematics of the Internet: E-Auctions and Markets*. Springer-Verlag, Berlin, 2001.

[8] Bikhchandani, S. and Mamer, J. Competitive Equilibrium in an Exchange Economy with Indivisibilities. *Journal of Economic Theory*, 74(2):385–413, 1997.

[9] Bikhchandani, S. and Ostroy, J.M. The Package Assignment Problem. Technical report, Anderson School of Management and Department of Economics, UCLA, 2000.

[10] Bourbeau, B., Crainic, T.G., Gendreau, M., and Robert, J. Design for Optimized Multilateral Multi-commodity Markets. *European Journal of Operations Research, Forthcoming*, 2004.

[11] Camerini, P.M., Fratta, L., and Maffioli, F. On improving relaxation methods by modified gradient techniques. *Mathematical Programming Study*, 3:26–34, 1975.

[12] Cheng, J.Q. and Wellman, M.P. The WALRAS Algorithm: A Convergent Distributed Implementation of General Equilibrium Outcomes. *Computational Economics*, 12:1–24, 1998.

[13] Dantzig, G.B. *Linear Programming and Extensions*. Princeton University Press, Princeton, 1963.

[14] de Vries, S. and Vohra, S. Combinatorial Auctions: A Survey. *INFORMS Journal on Computing*, 15(3):284–309, 2003.

[15] Demange, G., Gale, D., and Sotomayor, M. Multi-item auctions. *Journal of Political Economy*, 94:863–872, 1986.

[16] DeMartini, C., Kwasnica, A.M., Ledyard, J.O., and Porter, D.P. A New and Improved Design for Multi-object Iterative Auctions. Social Science Working Paper SSWP 1054, Caltech, 1999.

[17] Everett, H. Generalized Lagrange Multiplier Method for Solving Problems of Optimum Allocation of Resources. *Operations Research*, 11:399–417, 1963.

[18] Geoffrion, A.M. Elements of Large-Scale Mathematical Programming. *Management Science*, 16(11):652–685, 1970.

[19] Geoffrion, A.M. Duality in Nonlinear Programming: A Simplified Applications-Oriented Development. *SIAM Review*, 13(1):1–37, 1971.

[20] Gul, F. and Stacchetti, E. The English Auction with Differentiated Commodities. *Journal of Economic Theory*, 92:66–95, 2000.

[21] Hiriart-Urruty, J. and Lemaréchal, C. *Convex Analysis and Minimization Algorithms II: Advanced Theory and Bundle Methods.* A Series of Comprehensive Studies in Mathematics. Springer, New York, 1991.

[22] Kelley, J.E. The cutting-plane method for solving convex programs. *J. SIAM*, 8:703–712, 1960.

[23] Kim, K.H., Chung, W.J., Huang, H., and Chang, S.K. An Electronic Brokering Process for Truckload Freight. Technical report, Department of Industrial Engineer, Pusan National University, 2002.

[24] Kiwiel, K.C. Proximity Control in Bundle Methods for Convex Nondifferentiable Optimization. *Mathematical Programming*, 46:105–122, 1990.

[25] Kurose, J.F. and Simha, R. Microeconomic Optimal Resource Allocation in Distributed Computer Systems. *IEEE Transactions on Computers*, 38(5):705–717, 1989.

[26] Kutanoglu, E. and Wu, D. On Combinatorial Auction and Lagrangian Relaxation for Distributed Resource Scheduling. *IEEE Transactions*, 31:813–826, 1999.

[27] Larrson, T., Patriksson, M., and Strömberg, A-B. Ergodic Primal Convergence in Dual Subgradient Schemes for Convex Programming. *Mathematical Programming*, 86:283–312, 1999.

[28] Lasdon, L.S. *Optimization Theory for Large Systems.* Macmillan, London, 1970.

[29] Lemaréchal, C. Nondifferentiable Optimization. In Nemhauser, G.L., Rinnoy Kan, A.H.G, and Todd, M.J., editors, *Handbooks in Operations Research and Management Science, Vol.1.* North-Holland, Amsterdam, 1989.

[30] Malone, T.W., Yates, J., and Benjamin, R.I. Electronic Markets and Electronic Hierarchies. *Communications of the ACM*, 30(6):484–497, 1987.

[31] Mas-Colell, A., Whinston, M.D., and Green, J.R. *Microeconomic Theory.* Oxford University Press, 1995.

[32] Nisan, N. The Communication Complexity of Combinatorial Auctions. Draft, Institute of Computer Science, The Hebrew University of Jerusalem, 2001. Available at: http://www.cs.huji.ac.il/~noam/mkts.html.

[33] Parkes, D.C. iBundle: An Efficient Ascending Price Bundle Auction. In *Procedings of the ACM Conference on Electronic Commerce*, pages 148–157. ACM 2000, Minneapolis, MN, 2000.

[34] Parkes, D.C. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency.* PhD thesis, University of Pennsylvania, 2001.

[35] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, pages 295–308, Hidden Valley, Pennsylvania, 1993.

[36] Schramm, H. and Zowe, J. A Version of the Bundle Idea for Minimizing a Nonsmooth Function: Conceptual Idea, Convergence Analysis, Numerical Results. *SIAM Journal on Optimization*, 2(1):121–152, 1992.

[37] Tan, J.C. and Harker, P.T. Designing workflow coordination: centralized versus market-based mechanisms. *Information Systems Research*, 10:328–342, 1999.

[38] Walras, L. *Eléments d'économie politique pure*. L. Corbaz, Lausanne, 1874.

[39] M.P. Wellman. Market-oriented programming: Some early lessons. In Scott Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, River Edge, New Jersey, 1996.

[40] Wellman, M.P. A market-oriented programming environment and its application to multi-commodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–22, 1993.

[41] Wellman, M.P., Walsh, W.E., Wurman, P.R., and MacKie-Mason, J.K. Auction Protocols for Decentralized Scheduling. *Games and Economic Behavior*, 35:271–303, 2001.

[42] Wolfe, P. A Method of Conjugate Subgradients for Minimizing Nondifferentiable Convex Functions. *Mathematical Programming Study*, 3:145–173, 1975.

[43] Yamaki, H., Wellman, M.P., and Ishida, T. A market-based approach to allocating QoS for multimedia applications. In *Second International Conference on Multi-Agent Systems*, pages 385–392, Kyoto, 1996.

[44] Ygge, F. *Market-Oriented Programming and its Application to Power Load Management*. PhD thesis, Department of Computer Science, Lund University, 1998.

[45] Ygge, F. and Akkermans, H. Resource-Oriented Multi-Commodity Market Algorithms. *Autonomous Agents and Multi-Agent Systems Journal*, 3(1):53–71, 2000.